The Buyer's Guide to RAID Storage

StorageWorks

PERSPECTIVES

It's a busy Saturday, and one of the biggest order days of the year in your business. Transaction volume could hit a record. The work you've put into tuning the network (and the database on the server) is paying off. But three hours before close of business, disaster strikes. A month ago, you put an expansion chassis on your main server, adding another five gigabytes of storage. Now the whole bank of expansion drives connected to the server has gone down!

You hear groans from customer service, and your phone rings. "Not to worry," you tell the agitated voice on the line, "We've prepared for the worst. We'll be back up in no time."

Two-and-a-half hours later, you've replaced the power supply in the expansion chassis (you had a spare) and rebooted the server. Then you discovered that a disk crashed when the power supply failed. You replaced the disk and restored from tape (with the usual frustrations in identifying, locating, and reading the backup tape). By then, the day was over, the order takers had gone home, and the sales they should have made in those three hours were gone. Forever.

Monday morning, you're explaining the unforeseeable to the operations manager. "We'd done everything right: mirrored system drives, nightly backups, brand name products, UPS. It was a fluke, and no one could have planned for it."

She isn't buying. "I talked with my neighbor Saturday. He says if we used RAID, we wouldn't lose any up-time. You did a good job with downsizing and all the data integrity and backup issues. But now the critical issue is data availability. We are out of business every minute the data is unavailable. The customer goes somewhere else, and we never get those sales back. If you can't figure out how to protect our data, I'll get someone who can. Is that clear?"

In this new era of downsizing and distributed applications, the old "glass house" disciplines that ensured data availability through intensive systems management are no longer economically feasible. But the need for that level of data availability is greater than ever. The answer is to automate data availability, just as we have automated so many other aspects of system management. And that is where RAID storage has an important impact.

What does RAID mean?

RAID means Redundant Array of Independent Disks. It is a way of configuring multiple disk drives to achieve high data availability. In many cases, RAID can deliver improved performance, as well. All implementations of RAID have one important thing in common: a RAID array, whether it contains two disk drives, or five, or twenty, looks like one or more large disk drives to the user. You use a RAID drive just like you would any other drive. You can partition it if you want, and no application changes are needed to realize the benefits of RAID. All you see is excellent storage availability and better performance.

Where did RAID come from?

The RAID concept was developed by a team of researchers at the University of California at Berkeley in 1987. The researchers were looking for a way to use small-capacity, inexpensive, PC-type disk drives as an alternative to the expensive, large-capacity, 14-inch drives then common on mainframe computers. RAID was the result.

As RAID has moved from concept to product, and as drive technology has evolved, the reason for using RAID has changed. In 1987, big drives cost significantly more than small drives, and used significantly more space, power, and cooling. Now that large-capacity 5.25-inch drives and 3.5-inch drives are becoming the standard for all systems, the price and performance distinctions between large and small drives have disappeared (along with the large drives). Today, RAID's primary role is as an automated way to enhance data availability, not as an ingenious cost cutting strategy.

Why is RAID important?

The reason is simple: RAID helps you achieve data availability levels that in the past were only possible on costly mainframe systems. In addition, with RAID you can tune your storage subsystem performance to your business and application needs, building upon and enhancing your entire network investment.

Popular Raid Levels Compared

RAID 0 (striping)

+ eliminates I/O "hot spots"
+ low cost per megabyte
- worst data availability

RAID 1 (mirroring)

+ excellent data availability
+ excellent performance after disk failure
+ twice as fast on reads as a single disk
- doubles drive costs

RAID 0+1 (striping plus mirroring)

+ excellent data availability
+ excellent performance after disk failure
+ twice as fast on reads as a single disk
+ eliminates I/O "hot spots"
- doubles drive costs

RAID 3 (striped data, one parity disk)

+ excellent data availability
+ highest data transfer rate, best for large data transfers
- I/O request rate no faster than a single disk

RAID 5 (striped data/striped parity)

+ excellent data availability
+ eliminates I/O "hot spots"
+ good for many small I/Os
+ more cost-effective than mirroring for larger capacities
- variable performance after disk failure
- slower writes than a single disk

RAID levels are not a quality ranking.

RAID is not a single, discrete product. RAID comes in several varieties, or levels. Each RAID level offers tradeoffs among availability, performance, and cost. The different levels of RAID are numbered 0 through 6 (e.g., RAID 1, RAID 5). The numbers are not a quality or performance ranking. They are simply a means of distinguishing among RAID levels. The most popular RAID levels for network servers are levels 1 (mirroring) and 5 (striped data/striped parity), with RAID 3 (striped data, one parity disk) a distant third.

Where does RAID fit in your total storage strategy?

RAID primarily offers dramatic increases in storage system availability. It also offers performance benefits, depending on application. In addition, RAID, because it allows multiple disks to be managed as a single disk, may help make life easier for systems managers.

RAID is not a solution for all I/O problems. If high performance is your primary goal, then semiconductor memory, in one or more forms, may offer greater performance benefits than RAID offers. Semiconductor memory includes main memory, different kinds of cache, and solid state disks.

Similarly, RAID does not replace bulk storage, such as magnetic tape and optical disk, that offers transportability and lower cost. Nor does RAID mean not backing up data: the fact is, most data loss is caused by human error, not storage system failures, so thorough backup processes are still required.

Digital™ offers RAID products for multiple environments, including Novell® NetWare,® SunOS,™ Solaris,® SCO™ UNIX,® and MS-DOS,® as well as OpenVMS.™ Our objective is to help you understand the advantages and disadvantages of different approaches to RAID and choose the right one for your applications.

How do you know if you really need RAID?

There are other ways to improve performance that are as effective or more effective than RAID. And you will meet your data integrity goals with programs for backup, security, and disaster recovery. But RAID is the best and probably only way to dramatically improve availability.

If availability is really of primary concern, if anything less than 100% availability means lost revenue, then the premium you might pay for RAID will be well worth the investment, and you should not hesitate to move to RAID. If the need isn't that clear and you have to cost justify RAID, you need to look at the cost of acquiring the RAID, versus the cost of the loss of availability.

The cost of lost availability

You can roughly estimate the cost of lost availability. Start by looking at lost productivity. If you have 200 nodes on a network, supporting workers whose fully-burdened cost is $20 per hour, and who spend 80% of their time on the system, you'll have a loss of approximately 200 x $20 x 80% = $3,200 per hour.

If you have 100 workstations, supporting engineers whose fully burdened cost is $70 per hour, and who spend 50% of their time on the network, the rate of loss is 100 x $70 x 50% = $3,500 per hour.

You can look at lost business, too. If the workers in the first example are order takers, and they usually average $400 per hour in sales, the impact is 200 x $400 = $80,000 per hour! Obviously, in that situation you would never allow a single point of failure. But even if you lost access to only 10% of your customer

It's just an average day in mid-week a few months later. Since the crash, you've added a 4 gigabyte RAID subsystem, attached to your server. It's a RAID 5 with redundant controllers, power supplies, and fans. The installation went fine, and you're confident that you've made a real step ahead in availability. But you haven't had a test—yet.

A little after noon, you see a message come up on your system monitor. You check your RAID subsystem, and its LEDs show that one of the RAID drives has gone into a warning mode. You use the utility software to put that drive into "failed" mode, and then you pull the failed drive out of the enclosure. You can't help feeling a little nervous, even though you've tested this a dozen times. As usual, the system doesn't stumble a bit.

There is no apparent physical damage to the drive or its connectors. So you go to your spares locker and pull out a brand new drive, confirm that it's the same type as the failed drive, and plug it into the empty slot. Within a few seconds, its status LEDs indicate it is operating, and your administrator's console tells you that the RAID array is reconstructing.

You walk across the hall into customer service. Things look pretty normal, and you're not going to tell anyone otherwise. You sound completely uninterested when you ask "how's it going?" but nobody really pays any attention. They're not even aware there's been a disk failure and the RAID array is still reconstructing. Within twenty minutes reconstruction is complete. And you're convinced that RAID is the right solution.

files, you're still looking at a sales loss of $8,000 per hour!

Even if you're not losing sales directly, loss of availability can slow down other customer-related activities, like field service, order tracking, or customer account information. Losses of "good will" and "customer satisfaction" are hard to define, but you know they eventually lead to lost sales.

The RAID premium

Look at RAID acquisition cost compared to the same amount of storage in a non-RAID configuration. You'll pay some premiums for RAID. If you choose a hardware-based, free-standing RAID subsystem (i.e., one that is not integrated into your primary network server), you will pay for packaging and power, a RAID controller, with its complex firmware, and a host adapter, all in addition to the drives you'll need to meet user capacity requirements.

You also pay a premium for drive redundancy. How much more you pay depends on what type of redundancy you want. With RAID 1 (mirroring), this premium is 100%. With RAID 5, the redundancy premium is only one drive for the entire array; as you add more RAID storage, this premium becomes relatively smaller. However, the price of RAID hardware today means that you should have a requirement for at least several gigabytes of storage before a RAID 5 subsystem is more cost effective than simple mirroring.

So as you look at the cost of RAID versus the same amount of non-RAID storage, you will get a pretty good idea of exactly how much extra you will pay for RAID. Now balance that against your cost of lost availability. How many one-hour or two-hour lost availability incidents will it take to pay for the RAID premium? If it's even close to the number you'd expect to have in a year, you should consider RAID seriously.

RAID and availability

Availability is usually the major benefit of RAID subsystems. In this context, availability is the ability of the system to continue accessing data despite a device failure.

Vendors often talk about "availability" as if it were a yes or no choice. Actually, from a systems perspective, availability is a continuum.The steps you can take to enhance availability range from backing up a hard drive once a week to maintaining disaster-tolerant systems with up-to-the-second data redundancy. For RAID, data redundancy is the key to high availability. There are two major types of data redundancy: mirroring and parity.

Data redundancy: mirroring

Mirroring is the simplest way to achieve data redundancy. (RAID level 1 is defined as mirroring.) For each data drive there is a second drive that contains exactly the same data. Data is written to both drives. Thus, if one drive fails, the other drive can provide an exact copy of the lost data immediately. This process can be referred to either as "mirroring" or "duplexing." In mirroring, there is one disk controller, and data is written serially, first to one drive and then to the mirror. This imposes a performance penalty. In duplexing, there are two controllers, and data is written to both drives in parallel, eliminating the performance penalty.

The data availability of mirrored drives is excellent, but at a cost. The cost is the need to purchase twice the data capacity your application requires. Even so, for small capacities—up to 2 or 3 gigabytes—mirroring may be the most cost effective way to achieve high availability.

Data redundancy: parity

The second way to provide data redundancy is to apply sophisticated mathematical coding techniques to produce parity data. (RAID levels 3, 5, and 6 use parity data.) Stored in the RAID array, parity data makes it possible to maintain access to data even if the physical drive on which the data was stored has failed. When there is a request to read data from the failed drive, the RAID automatically recreates the requested data from parity data stored on the other drives in the array.

The advantage of parity is that it requires significantly less extra disk space than mirroring, because the parity data for the whole array usually requires just one extra drive.

Shadow redundancy (RAID level 1)

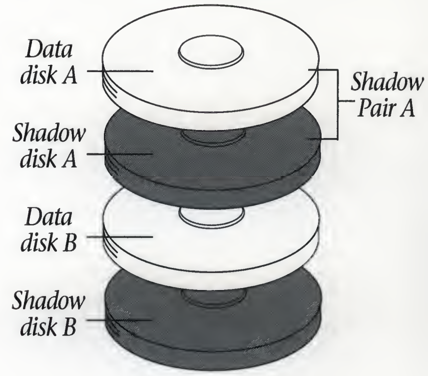

Parity redundancy (dedicated parity disk) (RAID level 3)

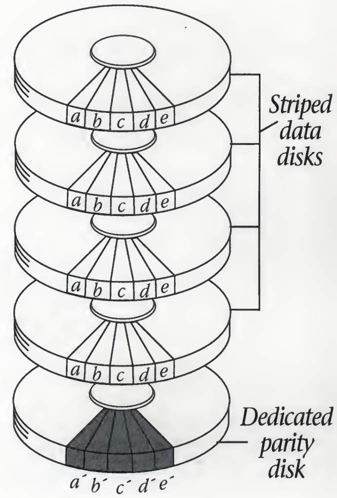

Parity redundancy (Striped parity/striped data) (RAID level 5)

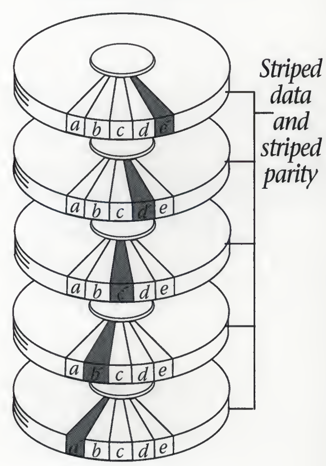

Hardware Redundancy Options

RAID installed in server

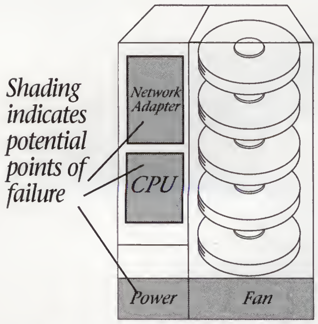

This configuration protects against disk drive failures only.

RAID attached to server, Redundant power supplies and fans

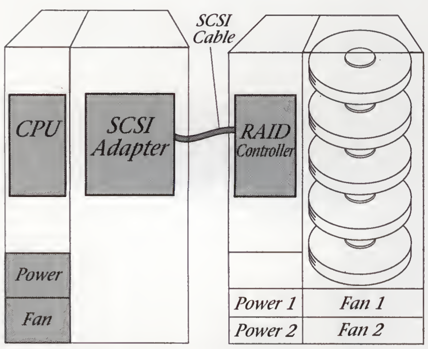

This configuration protects against failures in disks, array power supply, and array cooling.

RAID attached to server, for redundant hot-pluggable active components

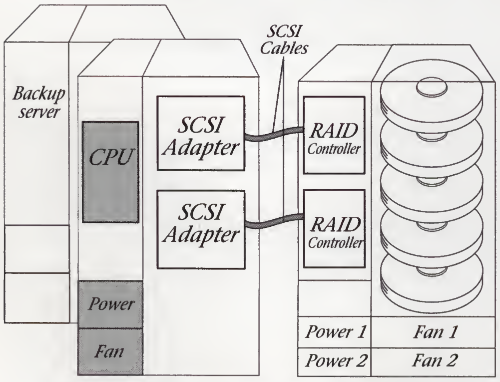

This configuration protects against failures in disks, array power supply, array cooling, RAID controllers, SCSI adapters and cables, and, optionally, CPU. Potential points of failure are passive components only.

Reconstruction

Data redundancy, whether it's based on mirroring or parity data, lets you maintain active access to your data in spite of a drive failure. Reconstruction is the process of rebuilding the data that was on the failed drive onto a replacement drive.

During reconstruction, there will be a fall-off in system performance, because controller resources and I/Os are being used to copy (RAID 1) or reconstruct data from parity (RAID 3,5,6). The speed of reconstruction depends on the available resources. This can be adjusted by the system manager. If transaction volume is high, he or she can slow reconstruction and devote more resources to transactions.

Most hardware RAID subsystems provide automated, on-the-fly reconstruction. However, with software RAID and with some older RAID products you may not get automated reconstruction, so you may need to shut down the system to reconstruct.

What about all the other things that can fail?

RAID primarily addresses availability at the drive level. However, drive failure is not the only way data access may be lost. CPU failures, power failures, cooling failures, controller failures, and cabling problems could also cut you off from your data.

If availability is an important goal for you, then your RAID subsystem should include redundant power, redundant cooling, dual data paths, dual controllers, an uninterruptable power supply (UPS) for the system server, and a good backup policy. Don't forget: RAID does not protect your data against accidental user deletion, application bugs, viruses, or natural disasters such as fires, floods, and earthquakes. So data on RAID subsystems still needs to be backed up regularly.

Most RAID subsystems allow "hot pluggable" spares, so you can pull out and replace a failed drive without shutting down your system. This makes possible the highest level of availability. Look for a RAID subsytem that has LEDs to tell a non-technical user which drive to replace, and that lets you replace a drive without a service call or even tools. With some RAID subsystems you can also "hot plug" other components, such as power supplies.

RAID and performance

RAID can't solve all performance problems. In fact, RAID is not primarily a performance enhancement. Having said that, if you want performance from RAID, which is best?

- RAID 0 offers excellent performance, but no availability benefit.
- RAID 1 offers good request-rate performance and high availability, with a big cost premium.
- RAID 0+1 offers the highest performance and high availability, but at a cost premium.
- RAID 3 offers the request-rate performance of a single drive, with significantly higher data rates and high availability.
- RAID 5 offers high request-rate performance, with better performance on reads than on writes, coupled with high availability.
- RAID 6 offers the highest availability and high request-rate performance, but imposes the most severe penalty in write performance.

As you can see, each RAID level has different performance characteristics. Why is this so, and what do these differences mean?

RAID and I/O Performance: Request rate, data transfer rate, read/write ratio, and "hot spots"

The impact of I/O workload on system performance depends on four basic elements: request rate, data transfer rate, read/write ratio, and hot spots (also known as "locality of reference").

- Request rate is the number of I/O requests per second the I/O subsystem is handling.
- Data transfer rate is how much user data is transferred per second by the I/O subsystem.
- Read/write ratio is the ratio of read requests to write requests: the mix between requests that merely copy data, and those that change it.
- Hot spots refers to the tendency of the workload to access data items located in close proximity to each other on a storage device.

A year later, availability has become an even more important issue in your business. Your organization is committed to offering the best customer support in the industry. As a result, you've gone to a "7 x 24" customer service schedule. So the pressure is onto maintain 100% system availability. You have confidence in your RAID approach, but now you need to go the extra mile. With redundant controllers, power supplies, and cooling, the RAID box itself is already optimized for availability. The most obvious point of failure now would be the server itself. Rather than go to the expense of a fully configured spare server, you upgrade one of the client PCs in customer service and designate it as the backup server. You add memory, load a backup copy of the operating system and the critical applications, and install a host adapter and cabling from the RAID box. With the fully redundant system, even if the server goes down, you can have the network back up and running in minutes. Of course, it hasn't happened yet. But now you're ready if it does.

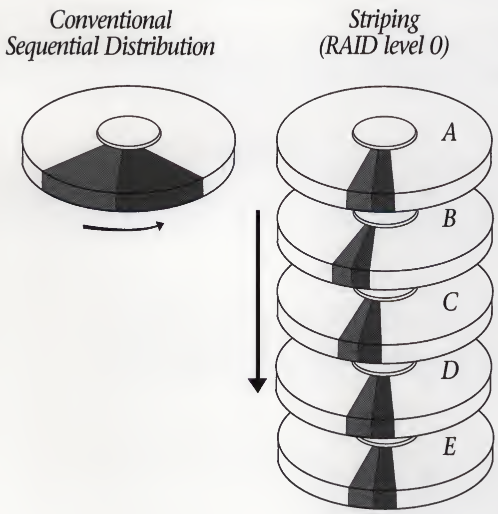

For most applications, request rate is the limiting I/O factor. General office automation, databases, transaction processing, and server applications tend to be request-rate intensive. In these applications, the number of requests is high, and the average request size tends to be small (under 8KB per request).

Some applications are limited by data transfer rate. Examples are CAD, imaging, and graphics. For these applications, request rate tends to be relatively low, but the average request size tends to be large (64KB and up). So the data rate requirement can be higher than a single drive can sustain, and RAID with multiple drives working together can offer a big performance benefit.

Read/write ratio is important in evaluating RAID levels, because some RAID levels perform much better reading than writing.

Hot spots are largely a problem with high request rate applications, since it is common for many requests to go to the same disk and overload that disk's request rate capacity.

Striping and performance (RAID 0)

The major source of RAID's performance benefit is striping of data. How does this work?

Striping enhances performance by spreading data across multiple drives (the "stripe set"). Striping works by first breaking user data into segments called "chunks". In a five-drive stripe set (whose members may be called A, B, C, D, E), the first chunk is placed on drive A, the second on B, the third on C, the fourth on D, the fifth chunk on E, the sixth back on A, and so forth until all data is stored. Striping is also known as RAID 0.

The system administrator can set chunk size based on application requirements. The relationship

between chunk size and average request size determines whether striping maximizes performance for request rate or data transfer rate. If chunk size is set larger than the average request size, all of the drives may be able to service different I/O requests simultaneously, significantly increasing request-rate performance.

If chunk size is set smaller than the average request size, then multiple drives in a stripe set can participate in a single request in parallel, thereby increasing data transfer rate. This is most beneficial with large request sizes, for which data transfer time is a significant portion of total data access time.

Load balancing is an automatic outcome of striping. Without striping, frequently accessed data may be concentrated on a single drive, creating a hot spot, and that drive can become a bottleneck. Striping spreads that data among several drives, so the I/O workload is balanced across several drives, and total system performance benefits. This load balancing effect of striping accounts for much of the observed performance benefit of RAID subsystems.

Read/write ratio and performance

The read/write ratio characteristic of the workload is an important factor in choosing a RAID level, because the read performance of some RAID levels (1, 5 and 6 especially) is considerably better than the write performance. So the read percentage can affect the response time you see.

Both Novell and UNIX implement I/O caching, which reduces the total number of I/Os going to disk. This improves overall response times, but it also increases the percentage of writes. Under moderate I/O loads, this has little impact. However,

Chunk size large vs request size

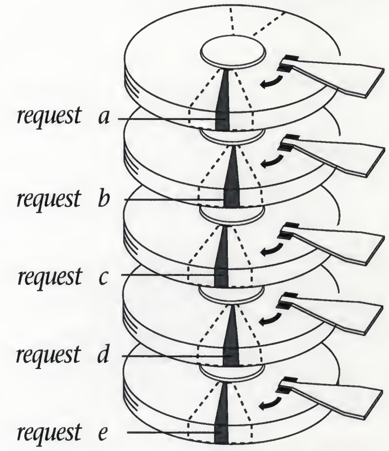

One actuator handles each request, many small request are handled simultaneously, and throughout improves (commerical, transaction-oriented applications).

Chunk size small vs request size

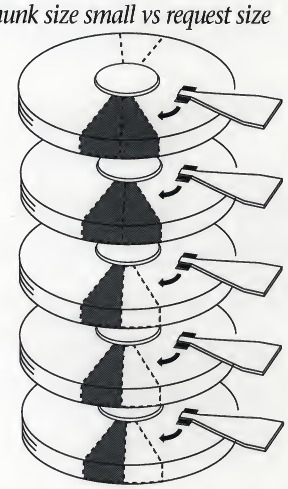

Multiple actuators simultaneously handle each large request, and bandwidth improves (technical applications involving large files).

Conventional disks (data distributed sequentially) *Striped array*

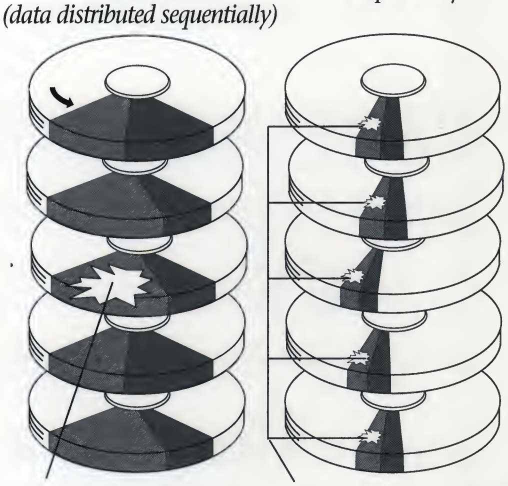

"Hot Spot"; 80/20 effect. Entire system is bottlenecked by the slowest data access. *"Hot Spot"is spread across multiple disks, improving overall performance.*

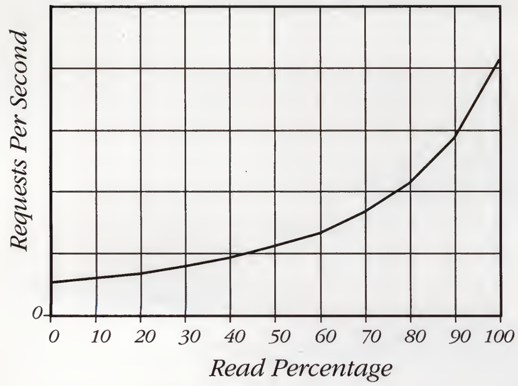

An example of the impact of read percentage on request rate performance.

when the total I/O activity gets close to saturating the RAID system, a high percentage of writes can degrade total I/O performance. The fall-off in performance under load can be sudden, not gradual, so make sure you benchmark under loads approximating real-world requirements. If consistent performance is important in your application, then your RAID subsystem should be sized to meet peak requirements, rather than average requirements.

Host-based or controller-based RAID?

RAID storage subsystems can be implemented in dedicated hardware or in software. In a hardware implementation, the RAID algorithms are packaged in the controller board that sits in your RAID subsystem, attached to the server I/O bus. In a software implementation, the RAID algorithms are incorporated in software that executes on your server CPU in concert with the operating system.

Controller-based solutions promise the highest performance over the widest range of application loads. This is true because, with host-based software solutions, the overhead of the RAID software increases as the load on the system increases. So the application software and the RAID software end up competing for host CPU cycles. This is especially true after a disk failure, when the original data must be reconstructed. With host-based software RAID, maximum request rates during reconstruction will be significantly lower, and reconstruction may take many hours. Dedicated controller-based RAID keeps reconstruction overhead off the CPU, minimizing the impact of drive failure on server operations. The possibility that the host-based, software solution will cost less is a compelling factor only if your system loads are consistently light, with few or no peak periods.

Software-based RAID can also leave you exposed in the event of a server failure; if the server goes down, your RAID system is down, too. Your RAID may be left in an ambiguous state, and at the very least some data will be lost. With hardware RAID, if the CPU fails, you can plug your RAID unit into a designated alternate server (it will need a host adapter card), and you won't miss a beat.

Summary of RAID benefits

Availability

RAID's greatest single benefit is availability. Remember, however, RAID only protects data at the drive level. For true high availability, RAID subsystems need to be engineered from the beginning for high availability. This means high-availability features such as dual paths to redundant power and cooling, dual paths to drives, redundant controllers, and "hot pluggable" spares capabilities.

Performance

RAID may offer performance benefits, depending on the RAID level. The primary performance benefit is the ability to increase either request rate or data transfer rate.

RAID also increases the potential variability of I/O subsystem performance. For example, if your RAID 3 subsystem (which is geared to handling relatively few large requests) is suddenly asked to handle many small requests, performance can degrade dramatically, bogging down your entire system. You should make sure your RAID subsystem gives you the flexibility to change RAID levels if your I/O profile changes.

RAID also offers the potential for increased read performance. However, workloads with a high percentage of writes may experience a performance penalty with some RAID levels. While this penalty may be reduced by cache techniques, it is still an issue.

Cost

RAID costs more than conventional drives of equivalent capacity, and some RAID levels cost more than others. Don't assume the high availability of RAID means savings on backup, either. You will need to back up your data, even with RAID. The idea behind RAID is to maximize availability and performance. By carefully defining your network characteristics and priorities, and by working closely with a system integrator, you should be able to realize these objectives and find a cost- effective solution.

RAID level summary

RAID 0 This is striping. User data is broken into chunks, which are stored onto the stripe set. No data redundancy is provided. Depending on chunk size, RAID 0 can significantly improve either data rate or request rate. Data rate is improved by making the chunk very small compared to an average request. As a result, all RAID set members are involved in all I/O requests, and effective data rate increases. Similarly, if the chunk size is large compared with the average request size, each request goes to a different drive, and request rate increases. The loss of any drive of the RAID set results in the loss of all data on that set. Since there is no redundancy, there is no cost overhead associated with RAID 0.

RAID 1 This is mirroring, or shadowing. RAID 1 provides at least two complete sets of all user data. Each member of the RAID set is duplicated. This means that all data remains available even when one member fails. Since all data is duplicated, a high level of disaster tolerance can be achieved by locating sets of duplicates in separate facilities. RAID 1 may also provide measurable performance improvement on read requests, because the controller can read from whichever drive is closest to the requested data.

When a drive fails, the user sees minimal impact on application performance and no loss of data availability. When the failed drive is replaced, performance will degrade as the replaced drive's data is written onto the new drive.

RAID 1 offers the possibility of not losing all data if multiple drives fail. For example, if there are four two-member mirror sets, and both members of one set fail, only 25% of data is lost.

The cost overhead for RAID 1 is 100%. To store four drives worth of data, you need eight drives.

RAID 0+1 This is the combination of striping and mirroring, implemented by striping mirror sets. RAID 0+1 provides the best performance of any type of RAID by combining the performance advantages of RAID 0 and RAID 1. It also provides disaster tolerance. Performance is better than RAID 0, and cost is the same as RAID 1.

RAID 2 is not now considered a practical solution, since more redundancy drives are required, and the performance is identical to that of RAID 3.

RAID 3 User data is striped across a set of drives. A drive to hold parity data is added. This data is calculated dynamically as user data is written to the other drives. Chunks are set small with respect to average request size. Typical chunk sizes are bits, bytes, or blocks. RAID 3 can improve data rate compared to a set of independent non-RAID drives. Request rate is limited to that of a single drive.

The cost overhead of RAID 3 is one drive, the parity drive. All data remains accessible even when any single drive fails.

RAID 4 is not now considered a practical solution, since all redundancy data is written to a single drive, causing a severe bottleneck. RAID 5 has the same cost and MTBF, with far superior performance.

RAID 5 Uses large chunk sizes and stripes each request and the parity data across all members of the RAID set. Since chunk size is large, usually only one drive participates in any request. This increases request rate performance, compared with a group of non-RAID drives. If a second drive fails before the first is replaced, access to all data in the RAID set is lost.

Since the chunks are large, insufficient data is received to enable parity to be calculated solely from the incoming data stream, as in RAID 3. The array controller must combine incoming data with existing parity data. Thus each write request involves reading from two drives (old data, old parity) and writing the new data onto two drives (new data, new parity). This results in poorer write performance than the other RAID levels.

The cost overhead of RAID 5 is one drive, the parity drive. All data remains accessible even when one drive fails. RAID 5 has considerably poorer performance than RAID 3 when a drive fails. If a second drive fails before the first is replaced, access to all data in the RAID set is lost.

RAID 6 RAID 6 offers very high availability. Two drives are used for redundancy, with sophisticated error-correcting codes. This enables RAID 6 to insure data integrity and availability even when two drives fail. Data and error-correcting information are striped across all members of the RAID set. Write performance is somewhat worse than RAID 5, because three drives must be accessed twice during writes. Request rate performance is high because chunks are large relative to average request size.

The cost overhead of RAID 6 is two drives. RAID 6 offers the poorest write performance and highest availability of any RAID level. Request rate performance for reads is high, comparable to groups of independent drives.

RAID level comparison

RAID Type	Description	Relative Availability	Request rate (Read/Write)	Data rate (Read/Write)	Cost Factor (1)	Types of Applications
Level 0	Striping; No redundancy	Proportionate to number of drives; worse than single drive	Based on chunk-size/request-size ratio. Can optimize for request rate or data rate Large chunks: Excellent	Small chunks: Excellent	1.0	Applications requiring high performance for non-critical data
Level 1	Shadowing. Both shadow-set members need to be written, degrading write performance.	Excellent	Good/Fair	Fair/Fair	2.0	System drives, critical files
Level 0+1	Striping plus shadowing together. Both shadow-set members need to be written, degrading write performance.	Excellent	Based on chunk-size/request-size ratio. Can optimize for request rate or data rate Large chunks: Excellent	Small chunks: Excellent/Good	2.0	Any critical response-time application
Level 3	Striped data with dedicated parity drive. Drives are rotationally synchronized.	Excellent	Poor	Excellent	1.25	Large I/O request size applications, such as imaging, CAD
Level 5	Striped data and parity.	Excellent	Excellent/Fair	Fair/Poor	1.25	High request rate, read-intensive, data lookup
Level 6	Striped data and parity with two parity drives.	Best	Excellent/Poor	Good/Poor	1.5	High request rate, read-intensive, data lookup
Individual drives	No RAID. No redundancy.	Proportional to number of drives.	Identical to single drive	Identical to single drive	1.0	

(1) **Cost factor** is the approximate multiplier of ordinary drive cost to achieve a given level of RAID. RAID 3, 5, and 6 require parity data, which means adding one drive (RAID 3, 5) or two drives (RAID 6) in addition to required user capacity. In addition to the parity drive, you may want to have a spare drive available to serve as a hot spare. This overhead is *not* counted in the table. The cost factor in this table assumes a 4-drive user capacity RAID set. A larger RAID set would change the cost factors for RAID 3, 5, and 6. The cost factor does not include the costs of power, packaging, or the RAID controller or software.

StorageWorks

PERSPECTIVES

StorageWorks Perspectives is a series of publications sponsored by Digital's Storage Business Unit. Each StorageWorks Perspective describes an important storage technology, user strategy, or product trend. Our goal is to give you a straightforward, informational presentation that will help you evaluate and select a storage strategy to fit your specific business needs.

StorageWorks is Digital's new line of industry-standard storage solutions that meet the current and future data storage needs of computer users in desktop to data center environments.

digital™